Photonics in the Package for Extreme Scalability (PIPES)
Frequently Asked Questions (FAQ) Document
Updated 11/20/2018

Q1: The BAA mentions "optical fiber" a number of times. Is using optical fibers a mandated approach?
A1: In TA1 and TA2, we anticipate that signals will be coupled to fiber in order to satisfy the link reach metric. For interfacing to the electronic IC core, the use of optical waveguides as an intermediate medium is allowed, provided optical coupling and propagation losses can be tolerated in the link budget. In TA3, there is no preference on the use of fiber vs. waveguides. The elimination of all fibers and waveguides (i.e., free space optical links) are not of interest.

Q2: Our organization's expertise is focused on a specific part or component of a photonic link, but does not cover a complete PIPES solution. How can we participate on the program?
A2: Proposals offering partial solutions will be considered incomplete as they will not address the full complement of program metrics and goals. Proposers are strongly encouraged to form multidisciplinary teams with broad expertise in order to contribute complete solutions to the respective technical areas.

Q3: Our organization is considering a proposal to TA1B with a specific application in mind. However, the metrics in TA1 do not correspond well to the application considered. Should we concern ourselves with the TA1 technical metrics?
A3: The TA1 metrics listed in Table 1 of the BAA do not apply to TA1B proposers, but indicate the expected direction of technology development. Proposers in TA1B should define and specify their own metrics based on the application proposed, while considering the technology development path outlined in the BAA. During the course of Phase 1 of the program, TA1B performers may have the opportunity to influence some characteristics of the technologies developed in TA1, and should identify potential means of using elements of PIPES technology for their specific application.

Q4: Will TA1 performers develop variants of their MCMs to satisfy specific TA1B performer requirements?
A4: No. TA1 performers will deliver broadly applicable packaged MCMs as outlined in the BAA. However, technology developed in TA1 such as IP blocks, I/O chiplets, foundry flows, packaging capabilities, and interface definitions will be available for use within the PIPES ecosystem. A stated goal of the PIPES program is to ensure that such developed technologies are accessible for custom DoD applications. However, extensive customization of MCM hardware lies outside the scope of the current BAA.

Q5: We have in mind a custom ASIC targeting a specific DoD application. Is it in scope to develop such ASIC under TA1 or TA1B of the program?
A5: No. The development of new ASICs may be possible within the TA1B Demonstration Phase, but such work is not to be proposed to the existing BAA. However, outlining the concept,

defining the characteristics, and describing the development effort of the project is encouraged.

Q6: Is the aggregate bandwidth defined in the BAA unidirectional or bidirectional?
A6: Optical communication will be bi-directional. For example, the 10T metric requires full duplex links with 10 Tb/s input to, and 10 Tb/s output from the package. This is discussed on pgs. 12 and 17 of the BAA, for TA1 and TA2 respectively.

Q7: Does the latency specified include error correction? Serialization/deserialization?
A7: Yes. The latency should include all signal conversions and signal processing required on both ends of the link. This includes overhead associated with error correction, equalization, SERDES, retiming, and any other signal delays caused by circuits within the link.

Q8: What should performers deliver to facilitate testing of the deliverables? Are they expected to deliver a full test bench?
A8: No. Laboratory equipment or test instrumentation is not required as part of the delivery. Any custom component or chip that was developed specifically as an aid for testing should be delivered. For example, an IC socket, a test board or a development board can be a part of the deliverable. Any external laser source with all required connections and couplers should also be delivered. Specifications of all electrical and optical interfaces as well as test procedures and methodologies should be provided and accompanied with suitable documentation.

Q9: Is there a 2" x 2" requirement for the maximum MCM size?
A9: The BAA refers to a preferred maximum size and form-factor for the package to maintain consistency with accepted microelectronics art. However, the preferred definitions are provided as guidance, and the specifications are not metrics of the program.

Q10: Is development of a new processing core in the scope of PIPES? Is modification of an existing core considered in scope?
A10: Development of a new processing core is not in the program scope. Modifications to existing designs to ensure the electrical interface between the core and the photonic I/O is in scope. Development of a "proxy core" with the sole function of demonstrating operation and providing a means to test I/O functionality and performance is considered in scope.

Q11: Is there a preference for using serial vs. parallel electrical interfaces?
A11: No.

Q12: Is there a preference for single mode vs. multi-mode fibers?
A12: No.

Q13: Is there a preference for using silicon vs. III-V photonic components?
A13: No.

Q14: Can a small company lead a proposal?
A14: Yes. There are no limitations on the size of proposing organizations, provided the proposing team demonstrates the capabilities needed to address the technical goals in their entirety.

Q15: Are there requirements for the size and power of a potential off-chip laser source?
A15: There are no requirements for size. The power dissipated should be fully accounted in the link budget.

Q16: Does the "proxy chip" power dissipation count in the energy efficiency calculation?
A16: The power required for data generation and test circuits does not count towards the link power consumption. However, the power required to drive any on-chip electrical interfaces to and from the photonic I/O should be included in the energy calculation.

Q17: Can one PI participate in multiple teams and proposals?
A17: Yes. If more than one proposal from a single performer is selected and duplicative work is identified, the work or costing may be modified during contract negotiations.

Q18: Is there a requirement for symmetric bi-directional operation of the photonic links?
A18: Yes. Both in TA1 and TA2, the links should demonstrate symmetric bi-directional operation. However, they do not have to connect identical ICs. For example, connecting an FPGA with an ASIC, GPU, or CPU is within scope. Connecting different ASICs is also in scope.

Q19: With respect to the TA1 Phase 2 deliverable of 10x MCM units, does the 1 pJ/b energy per bit (as defined in metrics) apply to the full EOE link? If so, the existing available fully functional IC cores (e.g. CPU, GPU, FPGA, ASIC) to be used for this demo will almost certainly have their own high speed SERDES which will not meet the energy per bit metric. Are we to count the on-chip I/O for these existing ICs in the energy per bit metric?
A19: Yes. The 1 pJ/bit metric applies to the full EOE link, and the on-chip I/O for the IC should be accounted for in the energy per bit metric. There is no assumption that the existing high speed SERDES will be used, and modifications to existing electrical interfaces, if necessary, is within scope of the PIPES program.

Q20: In order to achieve the MCM demo with the existing fully functional IC cores, i.e. using the "same I/O technology at 1 pJ/bit" (page 11 Phase 2), requires a new interface be built to convert between the existing fully functional IC core's electrical I/O format and the input to the optical I/O engine. Is that what is intended here? This interface would only be for the purposes of this MCM demo and serves no other useful purpose. Does this interface need to be included in the energy per bit total?
A20: A new interface to convert between the existing fully functional IC core's electrical I/O format and the input to the optical I/O engine is within scope, if required by the approach. This interface is to be included in the energy per bit total.

Q21: What aggregate bandwidth needs to be demonstrated with the fully functional MCM IC?
A21: The aggregate bandwidth for MCMs in TA1 Phase 2 is defined by the proposer. The overarching program goal is to develop in-package photonic I/O to enable efficient, high-bandwidth communications for future microelectronic systems, and aggressive targets consistent with observed scaling trends are anticipated here.

Q22: In reference to the TA1 Phase 2 "10T Technology Demonstrator", does this demonstrator need to show a full 10 Tbps, or just be scalable to this aggregate bandwidth? If the full 10 Tbps, does it need to show operation at the full aggregate bandwidth (i.e., all ports operating at once), or can the ports be operated sequentially?
A22: The demonstrator needs to show the full aggregate 10 Tb/s bandwidth. Ports may be tested sequentially, but should operate simultaneously.

Q23: In Table 1, TA1 Program Metrics, where is the operating temperature (to comply with the room T to 80C metric) to be measured? (e.g. At the heatsink surface? The heatsink/optical I/O "chip" interface? Inlet air temperature? Or somewhere else?)
A23: The intent is to demonstrate the photonic I/O can withstand operational temperatures that can be commonly found inside an IC package. Proposers are encouraged to describe a testing methodology that convincingly demonstrates the capability.

Q24: The TA2 Phase 3 deliverable on page 16 describes a demo with 100 Tbps aggregate bandwidth and 0.1 pJ/b energy, "consistent with 1 Pbps capability". This contradicts the milestone on page 22, which says 1 Pbps aggregate bandwidth at 0.1 pJ/b is required.
A24: The description of the deliverable on pg. 16, including in the text and in Table 2 is correct. DARPA will issue an amendment to clarify the wording of the milestone listed on pg. 22.

Q25: Are foreign entities (corporations) allowed to be participants of the proposal? If the answer is yes are there any constraints/requirements regarding the foreign participant? Are they allowed to receive program funds?
A25: Foreign participation is addressed on pg. 27 of the BAA, section III.A.1.C. Foreign participation is allowed to the extent that such participants comply with any necessary nondisclosure agreements, security regulations, export control laws, and other governing statutes applicable under the circumstances. In addition, potential proposers should self-assess their ability to address the goals of the Electronics Resurgence Initiative (ERI) and the PIPES program to create a domestic and trusted manufacturing and supply chain for the needs of the United States Department of Defense.

Q26: Can a university lead an effort?
A26: Yes, a university can lead an effort. However, as noted in the BAA, the anticipated funding type for TA1 and TA1B is 6.3 and, as such, a university prime contractor would need to be willing and able to accept that the work is NOT fundamental research (is restricted research - meaning, for example, Government pre-publication review and approval would be required). As a reminder, this same restriction applies for TA1 and TA1B whether a university is serving as the prime contractor or as a subcontractor.

Q27: How much R&D vs. commercialization is expected?
A27: PIPES assumes a balanced approach of research vs. commercial development, consistent with meeting the goals and metrics described in the BAA. On one hand, the technical goals exceed the current state-of-the-art by a large degree, thereby requiring a strong research element. On the other hand, commercial pull and acceptance of the technology by the broader microelectronics industry are recognized as essential for the ultimate success of the program. R&D approaches that have no path to manufacturing and commercialization will not be considered favorably.

Q28: Currently, no SoC has I/O energy below 1 pJ/bit, and the PIPES specifications has at least two such interfaces. It is hard to see how TA1 Phase 2 demo can be met without revision of the commercial IC.
A28: There is no assumption that the existing high-speed SERDES interfaces must be used, and modifications to the electrical interfaces of existing IC cores is within program scope.

Q29: Does the TA3 Reconfigurable Switching Thrust require us to deliver a fully packaged chip with ~2000 optical I/Os? Or can the prototype chip be demonstrated tested using bench top photonics testing such as a lensed fiber coupled to an inverse taper, or angled fiber coupler to a grating. Similarly, to provide the electrical signals for reconfiguration/switching, do we need to propose a companion ASIC with connectivity to the optical chip, or can we use electrical probes or a probe card?
A29: Delivery of fully packaged components, with fibers or waveguides and suitable electronic control, is required in Phase 3. As described in the BAA, demonstrator units should be provided with adequate instructions to support government testing and evaluation using standard laboratory equipment.

Q30: Who is the source selection authority?
A30: For Broad Agency Announcements (BAAs), DARPA does not have a Source Selection Authority (SSA). Instead, for BAAs, DARPA has what is referred to as a Scientific Review Official (SRO). The SRO is generally the Technical Officer Director.

Q31: In what program element (PE) does the PIPES program fall? Who is the PE program manager?
A31: The PIPES program funding uses program elements 0602716E and 0603739E. The DARPA Comptroller manages accounts for DARPA funding.

Q32: What technical readiness level (TRL) is required at the end of the program?
A32: TRL is not specified in the BAA. Rather, please refer to the quantitative metrics specified in the BAA.

Q33: Are TA1 and TA1B considered "not fundamental" only in consideration of publications, or also in terms of personnel, etc.?
A33: TA1 and TA1B will use "Advanced Technology Development" (6.3) funding, and contracted work will therefore be considered "not fundamental" in all aspects. Proposed efforts should be consistent with the requirements for 6.3 DoD funding.

Q34: Can a university participate in TA1?
A34: TA1 will use "Advanced Technology Development" (6.3) funding. Universities may be a performer in TA1, if willing and able to comply with non-fundamental research restrictions (applicable to both prime and subcontract participants).

Q35: Will DARPA provide a waiver stating that PIPES contracts and deliverables are not ITAR?
A35: No. DARPA does not make declarations of ITAR. Export control determination is the responsibility of the performer, working with the Department of State and Department of Commerce, as applicable.

Q36: For TA3, does the 1000 x 1000 switch metric require switching all wavelengths at the same time?
A36: The switching metric applies to optical ports as defined in the BAA. This implies that all wavelengths on a single port will be switched together.

Q37: Is cost sharing required?
A37: Cost sharing is highly encouraged for TA1 performers because there are clear commercial applications. Cost share will be favorably considered for all TAs.

Q38: Is there a standard wavelength for the photonics components.
A38: No.

Q39: Will Proposers Day slides be publicly available?
A39: Yes, PIPES Proposers Day slides will be posted on the DARPA/MTO Opportunities website.

Q40: Does the funding identified for each TA apply across all three phases of the program?
A40: Yes. The funding for each TA as indicated in the BAA includes all performers over all phases of the PIPES program.

Q41: Which elements of the program may be fabricated or performed outside of a domestic setting?
A41: The stated goal of the Electronics Resurgence Initiative is to create unique and differentiated domestic manufacturing capabilities accessible to the Department of Defense. As identified in the BAA, proposals should include a discussion describing how technology will be made accessible to the DoD via a trusted technology ecosystem. Detailed evaluation criteria are outlined in the BAA.

Q42: When is the estimated period of performance (POP) start date?
A42: The actual POP start date will depend on proposal selection and contracting. However, for purposes of cost proposal preparation, the BAA establishes an estimated period of performance start of July 2019 (use of 1 July 2019 is acceptable for these purposes).

Q43: Is there a minimum data rate?
A43: The BAA specifies aggregate I/O bandwidth, which represents the full duplex data rate of optical signaling from the package. Demonstrations must transmit and receive at this aggregate rate. Data rates per port, fiber, wavelength, channel, lane or other divisions of the total aggregate bandwidth are not specified and are at the proposer's discretion, commensurate with the proposed technical approach.

Q44: Should TA1B applications be focused on a problem today, or be focused on the 2025 to 2030 timeframe?
A44: Current DoD missions that are constrained by interconnect performance, as well as projected future needs, are both of interest. When projecting a need, proposers should provide sufficient evidence justifying the assumptions applied.

Q45: How does DARPA define "mission-relevant MCM" in the context of PIPES?
A45: For TA1, an MCM is defined as a multi-chip module that contains an advanced integrated circuit capability that could include CPU, GPU, FPGA, ASIC, or similar functionality. Proposals should address how the proposed MCM could positively impact applications of interest to the DoD.

Q46: Is there an upper bound on the technology node for proxy ICs?
A46: No.

Q47: Can you tell me if the references must also fit within the PIPES abstract 8-page and 5-page limits? It doesn't specifically call that out on page 31.
A47: Any references included in the Abstract will count towards the page count.

Q48: A question regarding the required performance of the optical switching fabric (Table 3). One requirement is that: "Switch reconfiguration time below 10 µs". Does the requirement apply to switching between two arbitrary states of the network? The context of this question is that for example in a cross-bar based network, a row-column type addressing (aka each row and each column is connected to a single electrical port) is possible w/o having to individually address all the switching elements at the cross points, but this comes at the expense that multiple switching operations have to be implemented to reconfigure the system between two arbitrary states. Is it correct to assume that in this case, the "switch reconfiguration time" alluded to in the FOA should include the time for all the operations needed to reconfigure the system between two arbitrary states?
A48: Yes, the reconfiguration time metric applies to switching between two arbitrary states.

Q49: A question regarding the requirement for "Switch power < 100 W" (Table 3). For a nonvolatile switching fabric, the average switching power has to be defined with respect to an intended reconfiguration frequency as maintaining a constant switch state does not consume any energy. What is the reconfiguration frequency that should be assumed in the proposal?
A49: The "Switch Power" metric in Table 3 of the BAA is intended as an upper bound for power consumption to ensure the technology is competitive with existing electronic solutions and consistent with packaging in traditional form factors. Allocation between static and dynamic power consumption is dependent on the technical approach considered. Proposals are expected to provide substantiated estimates on the power consumption of their specific implementation, and to note any assumptions or functionality restrictions (i.e., reconfiguration frequency limitations) needed to achieve the stated goals of the BAA.

Q50: Due to a natural disaster, our facility is closed and we have limited access to our systems. May the abstract due date be extended?
A50: The BAA has been amended (as of 16 November 2018) to include a revised Abstract Due Date of November 27th, 2018. Please see the BAA amendment on the Federal Business Opportunities website.

Q51: The TA3 section (p19) says, "Delivery of two operational demo units is required at the conclusion of Phase 3." However, Table 4 only calls for the Switching Thrust to deliver 2 packaged demonstrator units at the end of Phase 3. Is delivery of two demo units required at the end of Phase 3 for the TA3 Packaging Thrust as well?
A52: The delivery of two packaged demo units is only required in the TA3 Reconfigurable Optical Switching Technologies Thrust. Performers in the Packaging Thrust will be required to demonstrate fabrication capabilities and prototypes to meet program goals for integration and packaging, but demonstrator units are not a required deliverable.

Q52: I am interested obtaining in a participant list from Proposers Day, either with or without contact details. Could you provide the names of the individuals, or a list of the organizations/companies that attended?
A53: DARPA has not received permission from the participants of PIPES Proposers Day to distribute their names or contact information, and will not be providing an attendance list from Proposers Day.